B. Justification and Requirements for Continued Use of SUMEX

Computerized clinical data banks possess great potential as tools for assessing the efficacy of new diagnostic and therapeutic modalities, for monitoring the quality of health care delivery, and for support of basic medical research. Because of this potential, many clinical data banks have recently been developed throughout the United States. However, once the initial problems of data acquisition, storage, and retrieval have been dealt with, there remains a set of complex problems inherent in the task of accurately inferring medical knowledge from a collection of observations in patient records. These problems concern the complexity of disease and outcome definitions, the complexity of time relationships, potential biases in compared subsets, and missing and outlying data. The major problem of medical data banking is in the reliable inference of medical knowledge from primary observational data.

We see in the RADIX Project a method of solution to this problem through the utilization of knowledge engineering techniques from artificial intelligence. The RADIX Project, in providing this solution, will provide an important conceptual and technological link to a large community of medical research groups involved in the treatment and study of the chronic arthritides throughout the United States and Canada, who are presently using the ARAMIS Data Bank through the CIT facility via TELENET.

Beyond the arthritis centers which we have mentioned in this report, the TOD (Time-Oriented Data Base) User Group involves a broad range of university and community medical institutions involved in the treatment of cancer, stroke, cardiovascular disease, nephrologic disease, and others. Through the RADIX Project, the opportunity will be provided to foster national collaborations with these research groups and to provide a major arena in which to demonstrate the utility of artificial intelligence to clinical medicine.

C. Recommendations for Resource Development

The on-going acquisition of personal work-station Lisp processors is a very positive step, as these provide an excellent environment for program development, and can serve as a vehicle for providing programs to collaborators at other sites. Continued acquisitions are very desirable.

We also would hope that the central SUMEX facility, the DEC 2060, would continue to be supported. We continue to make constant use of this machine for text-editing, document preparation, file and database handling, communications, and program demos.

Responses to Questions Regarding Resource Future

Q: What do you think the role of the SUMEX-AIM resource should be for the period after 7/86, e.g., continue like it is, discontinue support of the central machine, act as a communications crossroads, develop software for user community workstations, etc.

A: In our opinion, the SUMEX 2060 should continue to be supported. The machine continues to be of value to us for text-editing (TV edit and emacs) and for document preparation (SCRIBE) and for communications and mail. We also depend on it as a central, reliable facility for program demos, for manipulating large databases, and maintaining central program files. It would be a real loss if it was discontinued.

Software for community work stations. Yes. Making good utility programs available to all users sounds like a good idea.

Q: Will you require continued access to the SUMEX-AIM 2060 and if so, for how long?

A: Yes. For the forseeable future and for the above reasons.

Q: What would be the effect of imposing fees for using SUMEX resources (computing and communications) if NIH were to require this?

A: We would pay them. The 2060 is worth it to us. Of course, if the fees were high, we would consider alternatives.

Q: Do you have plans to move your work to another machine workstation and if so, when and to what kind of system?

A: We are currently using two of the SUMEX Xerox 1108's for the development of our project. We will stay with these for the forseeable future.

IV.B. National AIM Projects

The following group of projects is formally approved for access to the AIM aliquot of the SUMEX-AIM resource. Their access is based on review by the AIM Advisory Group and approval by the AIM Executive Committee.

In addition to the progress reports presented here, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form.

IV.B.1. CADUCEUS Project

CADUCEUS Project

J. D. Myers, M.D. and Harry E. Pople, Jr., Ph.D.
University of Pittsburgh
Decision Systems Laboratory
Pittsburgh, Pa., 15261

I. SUMMARY OF RESEARCH PROGRAM

A. Project rationale

The principal objective of this project is the development of a high-level computer diagnostic program in the broad field of internal medicine as an aid in the solution of complex and complicated diagnostic problems. To be effective, the program must be capable of multiple diagnoses (related or independent) in a given patient.

A major achievement of this research undertaking has been the design of a program called INTERNIST-1, along with an extensive medical knowledge base. This program has been used over the past decade to analyze many hundreds of difficult diagnostic problems in the field of internal medicine. These problem cases have included cases published in medical journals (particularly Case Records of the Massachusetts General Hospital, in the New England Journal of Medicine), CPCs, and unusual problems of patients in our Medical Center. In most instances, but by no means all, INTERNIST-1 has performed at the level of the skilled internist, but the experience has high-lighted several areas for improvement.

B. Medical Relevance and Collaboration

The program inherently has direct and substantial medical relevance.

The institution of collaborative studies with other institutions has been deferred pending completion of the programs and knowledge base enhancements required for CADUCEUS. The installation of our own, dedicated VAX computer can be expected to aid considerably any future collaboration.

The INTERNIST-1 program has been used in recent years to develop patient management problems for the American College of Physician's Medical Knowledge Self-assessment Program, and to develop patient management problems and test cases for the Part III Examination and the developing computerized testing program of the National Board of Medical Examiners. In addition, selected other medical schools are employing the INTERNIST-1 knowledge base for medical student and house staff education.

----Accomplishments this past year

During 1983-84, under the supervision of Drs. Miller and Myers, Dr. Michael First, a former University of Pittsburgh medical student with extensive experience working in the Decision Systems Laboratory, developed a program called QUICK (QUick Index into Caduceus Knowledge), a prototypical electronic textbook of medicine utilizing the INTERNIST-1 knowledge base as its foundation. A paper describing QUICK, including an informal trial evaluating its utility, appears in the April 1985 issue of Computers and Biomedical Research. The residents in Internal Medicine who were given access to QUICK rated it favorably as a source of medical information. All three hospitals participating in the evaluation of QUICK have requested that they be given continued access to the program. An effort is being made to adapt QUICK to the IBM-PC for easier use by physicians.

From 1981 through 1983, Dr. Miller, under NLM New Investigator Award 5R23-LM03589, developed a clinical patient case simulator program, CPCS. The goal of the project was to build a program and knowledge base capable of constructing, de novo, logically consistent and clinically plausible artificial patient case summaries. Such a program would be useful in helping medical students to broaden their diagnostic skills. The program might also be used in generating cases for testing purposes, as this is now done manually by the National Board of Medical Examiners for their certification examinations. CPCS was a successful feasibility study; its performance has not yet been formally evaluated. Plans have been made to convert the entire INTERNIST-1 knowledge base into the format used by CPCS, and to add a better representation of time to the CPCS program and knowledge base.

Drs. Miller and Myers have developed, as part of the CPCS project, a new format for the internal medicine knowledge base. The specific details of this format have been described in previous progress reports. We have, in a period of three to four man-months, converted on paper the INTERNIST-1 knowledge base for liver diseases into the new format. This represents about one-sixth of the entire INTERNIST-1 knowledge base.

Dr. Miller has written an editor program to enter and maintain the new knowledge base, using Franz Lisp. At present, that editor program has been used to construct some 15-17 diagnoses from the INTERNIST-1 liver diseases. This includes creation of some 50-70 facets describing the underlying pathophysiology. A total of 200-300 findings have been entered into the new knowledge base, and because of their complexity, they correspond to 400-600 INTERNIST-1 style manifestations. During the past year, two fellows in Computer Medicine, Drs. Lynn Soffer and Fred Masarie, have converted all INTERNIST-1 findings into the new format required by CPCS.

Dr. Miller has also written, over the past year, a new diagnostic program which uses the information in the new knowledge base as a substrate for making diagnoses in internal medicine. The program's behavior is roughly comparable to that of INTERNIST-1 on similar cases in the limited problem domain currently available for testing. This remains an area of continued research activity.

In addition to the aforementioned work in internal medicine, Drs. Gordon Banks and John Vries have been working on the development of a neurological diagnostic component for CADUCEUS. Dr. Banks has developed a neuroanatomic database which contains spatial descriptors for nearly 1,000 neuroanatomic structures and contains information as to their blood supply and function. This database will allow anatomic localization of neurologic lesions. Some of this work for the peripheral nervous system has been done previously by students in our laboratory. The approach to the central nervous system has been to design a set of "symbolic coordinates". In constructing the neuroanatomic database, the human body, including the nervous system, is conceptually partitioned into a set of cubes (boxes). Attached to each cube LISP atom are lists of all of the anatomic structures that are completely and partially contained within the cube, as well as the blood supply to the region. This structure facilitates rapid retrieval of the location of a given anatomic structure as well as rapid localization of possible areas of involvement when there is evidence of dysfunction of one or more neural systems.

The hierarchical arrangement of the nested cubes ensures rapid convergence during searches, because if the sought object is not found in a parent cube, there is no need to search for it in any of the patient's children cubes. The addition of anatomic reasoning may allow parsimonious explanation of multiple manifestations arising from a single lesion, or allow the program to query the user regarding the presence of manifestations of involvement of areas that might be expected to be affected by whatever clinical state the program has under current consideration.

The neuroanatomic database has been successfully complemented on the VAX 11/780. Efforts are currently underway to implement the system on lower cost AI workstations such as the SUN and the PERQ.

Dr. Vries has continued to work on an image processing system based on "octree" encoding. Sean McLinden has developed an interface to the General Electric 9800 series CT scanner that permits direct input of data from the scanner to the octree system. The octree system output consists of 3 dimensional shaded images of CT objects at 1 mm resolution. Three dimensional images containing 2 million pixels can be scaled, translated, and rotated by the system in 30-60 seconds.

An interface to the neuroanatomic database has also been developed that maps the 27-ary tree representational scheme of the database into an octree representational scheme. This has been used to implement an interactive program that allows a user to generate a three dimensional image of the brain by logically ORing database objects.

A prototype system for the automated diagnosis of CT scans has also been implemented. The system uses the flavors package, and the RUP truth maintenance system to reason about the distribution of CT densities in quadtrees (2 dimensional representations) or octrees (3 dimensional representations). Such a system might ultimately provide CADUCEUS with direct access to the diagnostic information in neuro images.

The medical knowledge base has continued to grow both in the incorporation of new diseases and the modification of diseases already profiled so as to include recent advances in medical knowledge. Several dozen new diseases have been profiled during the past year and the pediatrics knowledge base has continued to grow.

----Research in progress

There are five major components to the continuation of this research project:

1. The enlargement, continued updating, refinement and testing of the extensive medical knowledge base required for the operation of INTERNIST-I.

2. The completion and implementation of the improved diagnostic consulting program, CADUCEUS, which has been designed to overcome certain performance problems identified during the past years of experience with the original INTERNIST-I program.

3. Institution of field trials of CADUCEUS on the clinical services in internal medicine at the Health Center of the University of Pittsburgh.

4. Expansion of the clinical field trials to other university health centers which have expressed interest in working with the system.

5. Adaptation of the diagnostic program and data base of CADUCEUS to subserve educational purposes and the evaluation of clinical performance and competence.

Current activity is devoted mainly to the first two of these, namely, the continued development of the medical knowledge base, and the implementation of the improved diagnostic consulting program.

D. List of relevant publications

1. First, M.B., Soffer, L.J., Miller, R.A.:*QUICK (Quick Index to Caduceus Knowledge): Using the Internist-1/Caduceus knowledge base as an an electronic textbook of medicine*. Comput. Biomed. Res. April 1985.

2. Miller, R.A.: *Internist-1/CADUCEUS: Problems Facing Expert Consultant Programs. Methods of Information in Medicine*. Schattauer, Stuttgart New York, Vol. 23, No. 1, January 1984, pp. 9-14.

3. Miller, R.A.: *A Computer-based Patient Case Simulator. Clinical Research*. 1984, 32:651A. (abstract).

4. Miller, R.A., Schaffer, K.F., Meisel, A.: *Ethical and legal issues related to the use of computer programs in clinical medicine*. Annals of Internal Medicine. 1985, 102:529-536.

5. Myers, J.D.: *Educating future physicians: Something old, Something new*. Ohio State Univ. Proceedings of Symposium, Medical Education in the 21st Century. (in press.)

6. Myers, J.D.: *The process of clinical diagnosis and its adaptation to the computer IN The Logic of Discovery and Diagnosis in Medicine*. University of Pittsburgh Series in the Philosophy and History of Science, Univ. of California Press (in press).

7. Pople, H.E.: *CADUCEUS: An Experimental Expert System for Medical Diagnosis. IN The AI Business*. Edited by Patrick H. Winston and Karen A. Prendergast. 1984, pp. 67-80.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A, B. Medical Collaborations and Program Dissemination Via SUMEX

CADUCEUS remains in a stage of research and development. As noted above, we are continuing to develop better computer programs to operate the diagnostic system, and the knowledge base cannot be used very effectively for collaborative purposes until it has reached a critical stage of completion. These factors have stifled collaboration via SUMEX up to this point and will continue to do so for the next year or two. In the meanwhile, through the SUMEX community there continues to be an exchange of information and states of progress. Such interactions particularly take place at the annual AIM Workshop.

C. Critique of Resource Management

SUMEX has been an excellent resource for the development of CADUCEUS. Our large program is handled efficiently, effectively and accurately. The staff at SUMEX have been uniformly supportive, cooperative, and innovative in connection with our project's needs.

III. RESEARCH PLANS

A. Project Goals and Plans

Continued effort to complete the medical knowledge base in internal medicine will be pursued including the incorporation of newly described diseases and new or altered medical information on "old" diseases. The latter two activities have proven to be more formidable than originally conceived. Profiles of added diseases plus other information is first incorporated into the medical knowledge base at SUMEX before being transferred into our newer information structures for CADUCEUS on the VAX. This sequence retains the operative capability of INTERNIST-1 as a computerized "textbook of medicine" for educational purposes.

B. Justification and Requirements for Continued SUMEX Use

Our use of SUMEX will obviously decline with the installation of our VAX and the use of personal work stations. Nevertheless, the excellent facilities of SUMEX are expected to be used for certain developmental work. It is intended for the present to keep INTERNIST-1 at SUMEX for comparative use as CADUCEUS is developed here.

Our best prediction is that our project will require continued access to the 2060 for the next two to three years and we consider such access essential to the future development of our knowledge base. After that time, our work can probably be accomplished on our VAX and personal work stations such as Symbolics. The imposition of fees for the use of SUMEX facilities would seem to involve unnecessary book-keeping and probably would detract from the use of SUMEX, which is currently so efficient and pleasant.

Our team hopes to remain as a component of the SUMEX community and to share experiences and developments.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

Our predictable needs in this area will be met by our dedicated VAX computer and newly acquired personal work stations.

D. Recommendations for Future Community and Resource Development

Whether a program like CADUCEUS, when mature, will be better operated from centralized, larger computers or from the developing self contained personal computers is difficult to predict. For the foreseeable future it would seem that centralized, advanced facilities like SUMEX will be important in further program development and refinement.

IV.B.2. CLIPR - Hierarchical Models of Human Cognition

Hierarchical Models of Human Cognition (CLIPR Project)

Walter Kintsch and Peter G. Polson
University of Colorado
Boulder, Colorado

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The two CLIPR projects have made progress during the last year. The prose comprehension project has completed one major project, and is designing a prose comprehension model that reflects state-of-the-art knowledge from psychology (van Dijk & Kintsch, 1983) and artificial intelligence. During the last three years, Polson, in collaboration with Dr. David Kieras of the University of Michigan, has continued work on a project studying the psychological factors underlying device complexity and the difficulties that nontechnically trained individuals have in learning to use devices like word processors. They have developed formal representations of a user's knowledge of how to operate a device and of the user-device interface (Kieras & Polson, in Press) and have completed several experiments evaluating their theory (Polson & Kieras, 1984, 1985).

B. Technical Goals

The CLIPR project consists of two subprojects. The first, the text comprehension project, is headed by Walter Kintsch and is a continuation of work on understanding of connected discourse that has been underway in Kintsch's laboratory for several years. The second, the device complexity project is headed by Peter Polson in collaboration with David Kieras of the University of Michigan. They are studying the learning and problem solving processes involved in the utilization of devices like word processors or complex computer controlled medical instruments (Kieras & Polson, in Press)

The goal of the prose comprehension project is to develop a computer system capable of the meaningful processing of prose. This work has been generally guided by the prose comprehension model discussed by van Dijk & Kintsch (1983), although our programming efforts have identified necessary clarifications and modifications in that model (Kintsch & Greeno, 1985; Fletcher, 1985; Walker & Kintsch, 1985; Young, 1985). In general, this research has emphasized the importance of knowledge and knowledge-based processes in comprehension. We hope to be able to merge the substantial artificial intelligence research on these systems with psychological interpretations of prose comprehension, resulting in a computational model that is also psychologically respectable.

The goal of the device complexity project is to develop explicit models of the user-device interaction. They model the device as a nested automata and the user as a production system. These models make explicit kinds of knowledge that are required to operate different kinds of devices and the processing loads imposed by different implementations of a device.

C. Medical Relevance and Collaboration

The text comprehension project impacts indirectly on medicine, as the medical profession is no stranger to the problems of the information glut. By adding to the research on how computer systems might understand and summarize texts, and determining ways by which the readability of texts can be improved, medicine can only be helped by research on how people understand prose. Development of a more thorough understanding of the various processes responsible for different types of learning problems in children and the corresponding development of a successful remediation strategy would also be facilitated by an explicit theory of the normal comprehension process.

The device complexity project has two primary goals: the development of a cognitive theory of user-device interaction in including learning and performance models, and the development of a theoretically driven design process that will optimize the relationships between device functionality and ease of learning and other performance factors (Polson & Kieras, 1983, 1984, 1985). The results of this project should be directly relevant to the design of complex, computer controlled medical equipment. They are currently using word processors to study user-device interactions, but principles underlying use of such devices should generalize to medical equipment.

Both the text comprehension project and the device complexity project involve the development of explicit models of complex cognitive processes; cognitive modeling is a stated goal of both SUMEX and research supported by NIMH.

Several other psychologists have either used or shown an interest in using an early version of the prose comprehension model, including Alan Lesgold of SUMEX's SCP project, who is exporting the system to the LRDC Vax. We have also worked with James Greeno -- another member of the SCP project -- on a project that will integrate this model with models of problem solving developed by Greeno and others at the University of California, Berkeley. Needless to say, all of this interaction has been greatly facilitated by the local and network-wide communication systems supported by SUMEX. The mail system, of course, has also enabled us to maintain professional contacts established at conferences and other meetings, and to share and discuss ideas with these contacts.

D. Progress Summary

The version of the prose comprehension model of 1978 (Kintsch & van Dijk, 1978), which originally was realized as a computer simulation by Miller & Kintsch (1980), has been extended in a major simulation program by Young (1985). Unlike the earlier program, Young includes macroprocessing in her model, and thereby greatly extends the usefulness of the program. It is expected that this program will be widely useful in studies of prose where a detailed theoretical analysis is desired.

The general theory has been reformulated and expanded in van Dijk & Kintsch (1983). This research report of book length presents a general framework for a comprehensive theory of discourse processing. It has been applied to an interesting special case, the question of how children understand and solve word arithmetic problems, by Kintsch & Greeno (1985). A simulation for this model, using INTERLISP, has been supplied in Fletcher (1985).

The device complexity project is in its third year. They have developed an explicit model for the knowledge structures involved in the user-device interaction, and they are developing simulation programs. Their preliminary theoretical results are described in Kieras & Polson (in Press). They have also completed several experiments evaluating the theory (Polson & Kieras, 1984, 1985) and have shown that number of productions predicts learning time and that number of cycles and working memory operations predicts execution time for a method.

E. List of Relevant Publications

1. Fletcher, R. C.: *Understanding and solving word arithmetic problems: A computer simulation.* Technical Report N0. 135, Institute of Cognitive Science, Colorado, 1984.

2. Kieras, D.E. and Polson, P.G.: *The formal analysis of user complexity.* Int. J. Man-Machine Studies, In Press.

3. Kintsch, W. and van Dijk, T.A.: *Toward a model of text comprehension and production.* Psychological Rev. 85:363-394, 1978.

4. Kintsch, W. and Greeno, J.G.:*Understanding and solving word arithmetic problems.* Psychological Review, 1985, 92, 109-129.

5. Miller, J.R. and Kintsch, W.: *Readability and recall of short prose passages: A theoretical analysis.* J. Experimental Psychology: Human Learning and Memory 6:335-354, 1980.

6. Polson, P.G. and Kieras, D.E.: *Theoretical foundations of a design process guide for the minimization of user complexity.* Working Paper No. 3, Project on User Complexity, Universities of Arizona and Colorado, June, 1983.

7. Polson, P.G. and Kieras, D.E.: *A formal description of users' knowledge of how to operate a device and user complexity.* Behavior Research Methods, Instrumentation, & Computers, 1984, 16, 249-255.

8. Polson, P.G. and Kieras, D.E.: *A quantitative model of the learning and performance of text editing knowledge.* Proceedings of the CHI 1985 Conference on Human Factors in Computing. San Francisco, April 1985.

9. Van Dijk, T.A. and Kintsch, W.:*STRATEGIES OF DISCOURSE COMPREHENSION.* Academic Press, New York, 1983.

10. Young, S.: *A theory and simulation of macrostructure.* Technical Report No. 134, Institute of Cognitive Science, Colorado, 1984.

11. Walker, H.W., Kintsch, W.: *Automatic and strategic aspects of knowledge retrieval.* Cognitive Science, 1985, 9, 261-283.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Sharing and Interactions with Other SUMEX-AIM Projects

Our primary interaction with the SUMEX community has been the work of the prose comprehension group with the AGE and UNITS projects at SUMEX. Feigenbaum and Nii have visited Colorado, and one of us (Miller) attended the AGE workshop at SUMEX. Both of these meetings have been very valuable in increasing our understanding of how our problems might best be solved by the various systems available at SUMEX. We also hope that our experiments with the AGE and UNITS packages have been helpful to the development of those projects.

We should also mention theoretical and experimental insights that we have received from Alan Lesgold and other members of the SUMEX SCP project. The initial comprehension model (Miller & Kintsch, 1980) has been used by Dr. Lesgold and other researchers at the University of Pittsburgh, as well as researchers at Carnegie-Mellon University, the University of Manitoba, Rockefeller University, and the University of Victoria.

B. Critique of Resource Management

The SUMEX-AIM resource is clearly suitable for the current and future needs of our project. We have found the staff of SUMEX to be cooperative and effective in dealing with special requirements and in responding to our questions. The facilities for communication on the ARPANET have also facilitated collaborative work with investigators throughout the country.

III. RESEARCH PLANS

A. Long Range Projects Goals and Plans

The goal of the prose comprehension project is to develop a computer system capable of the meaningful processing of prose. This work has been generally guided by the prose comprehension model discussed by van Dijk & Kintsch (1983), although our programming efforts have identified necessary clarifications and modifications in that model (Kintsch & Greeno, 1985; Fletcher, 1985; Walker & Kintsch, 1985; Young, 1985). In general, this research has emphasized the importance of knowledge and knowledge-based processes in comprehension. We hope to be able to merge the substantial artificial intelligence research on these systems with psychological interpretations of prose comprehension, resulting in a computational model that is also psychologically respectable.

The primary goal of the device complexity project is the development of a theory of the processes and knowledge structures that are involved in the performance of routine cognitive skills making use of devices like word processors. We plan to model the user-device interaction by representing the user's processes and knowledge as a production system and the device as a nested automata. We are also studying the role of mental models in learning how to use them.

B. Justification and Requirements for Continued SUMEX Use

Both the prose comprehension and the user-computer interaction projects have shifted their actual simulation work from SUMEX to systems at the University of Colorado and the University of Michigan. Both projects use Xerox 1108 systems continuing their work in INTERLISP. However, we consider our continued access to SUMEX critical for the successful continuation of these projects.

Access to SUMEX provides us with continued contact with the SUMEX community, which is especially critical for the prose comprehension project. Knowledge representation languages, e.g. UNITS, and other tools developed by SUMEX are critical for this project. Alternative sources of such software are typically unsatisfactory because the systems have only been developed for use on one project and are typically very poorly documented and less than completely debugged. We hope that our continued membership in the community will be offset by the input that we have been and will continue to provide to various projects: our relationship has been symbiotic, and we look forward to its continuation.

Access to SUMEX's mail facilities are critical for the continued success of these projects. These facilities provide us with the means to interact with colleagues at other universities. Kintsch is currently collaborating with James Greeno, who is at the University of California at Berkeley, and Polson's long-term collaborator, David Kieras, is at the University of Michigan. In addition, our access to the Xerox 1108 (Dandelion) user's community is through SUMEX.

We currently use four computing systems for the VAX 11/780, and three Xerox 1108s, one of which is at the University of Michigan. The VAX is used primarily to collect experimental data designed to evaluate the simulation models and to do necessary statistical analysis.

C. Needs and Plans for Other Computational Resources

SUMEX provides us with two critical needs. The first is communication, which we discussed in the preceding paragraph. The second is technical advice and access to various knowledge representation languages like UNITS.

We envisage our future needs to be communication currently served by the SUMEX 2060 and technical advice and necessary software provided by the SUMEX staff.

D. Recommendations for Future Community and Resource Development

Our future needs are for the SUMEX-AIM resource to act as a communications crossroad and to develop software and provide technical support for user community work stations. We have no preferences as to how such services are provided either with a communication server on the network or with the central machine like the current 2060.

We will continue to need access to the SUMEX-AIM 2060 in order to access communication networks and to interact with the SUMEX-AIM staff and community.

If communications and access to the staff are provided through some other mechanism, then we would no longer need access to the 2060.

We would be willing to pay fees for using SUMEX communication resources if required by NIH. However, our willingness is price sensitive. Any charges over $1,000 a year would mean we should communicate with people directly by long-distance telephone.

IV.B.3. MENTOR Project

MENTOR Project

Stuart M. Speedie, Ph.D.
School of Pharmacy
University of Maryland

Terrence F. Blaschke, M.D.
Department of Medicine
Division of Clinical Pharmacology
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The goal of the MENTOR (Medical EvaluatioN of Therapeutic ORders) project is to design and develop an expert system for monitoring drug therapy for hospitalized patients that will provide appropriate advice to physicians concerning the existence and management of adverse drug reactions. The computer as a record-keeping device is becoming increasingly common in hospital-based health care, but much of its potential remains unrealized. Furthermore, this information is provided to the physician in the form of raw data which is often difficult to interpret. The wealth of raw data may effectively hide important information about the patient from the physician. This is particularly true with respect to adverse reactions to drugs which can only be detected by simultaneous examinations of several different types of data including drug data, laboratory tests and clinical signs.

In order to detect and appropriately manage adverse drug reactions, sophisticated medical knowledge and problem solving is required. Expert systems offer the possibility of embedding this expertise in a computer system. Such a system could automatically gather the appropriate information from existing record-keeping systems and continually monitor for the occurrence of adverse drug reactions. Based on a knowledge base of relevant data, it could analyze incoming data and inform physicians when adverse reactions are likely to occur or when they have occurred. The MENTOR project is an attempt to explore the problems associated with the development and implementation of such a system and to implement a prototype of a drug monitoring system in a hospital setting.

B. Medical Relevance and Collaboration

A number of independent studies have confirmed that the incidence of adverse reactions to drugs in hospitalized patients is significant and that they are for the most part preventable. Moreover, such statistics do not include instances of suboptimal drug therapy which may result in increased costs, extended length-of-stay, or ineffective therapy. Data in these areas are sparse, though medical care evaluations carried out as part of hospital quality assurance programs suggest that suboptimal therapy is common.

Other computer systems have been developed to influence physician decision making by monitoring patient data and providing feedback. However, most of these systems suffer from a significant structural shortcoming. This shortcoming involves the evaluation rules that are used to generate feedback. In all cases, these criteria consist of discrete, independent rules. Yet, medical decision making is a complex process in which many factors are interrelated. Thus attempting to represent medical decision-making as a discrete set of independent rules, no matter how complex, is a task that can, at best, result in a first order approximation of the process. This places an inherent limitation on the quality of feedback that can be provided. As a consequence it is extremely difficult to develop feedback that explicitly takes into account all information available on the patient. One might speculate that the lack of widespread acceptance of such systems may be due to the fact that their recommendations are often rejected by physicians. These systems must be made more valid if they are to enjoy widespread acceptance among physicians.

The proposed MENTOR system is designed to address the significant problem of adverse drug reactions by means of a computer-based monitoring and feedback system to influence physician decision-making. It will employ principles of artificial intelligence to create a more valid system for evaluating therapeutic decision-making.

The work in the MENTOR project is intended to be a collaboration between Dr. Blaschke at Stanford and Dr. Speedie at the University of Maryland. Dr. Speedie provides the expertise in the area of artificial intelligence programming. Dr. Blaschke provides the medical expertise. The blend of previous experience, medical knowledge, computer science knowledge and evaluation design expertise they represent is vital to the successful completion of the activities in the MENTOR project.

C. Highlights of Research Progress

The MENTOR project was initiated in December 1983. The project has been funded by the National Center for Health Services Research since January 1, 1985. Initial effort has focused on exploration of the problem of designing the MENTOR system. Work has begun on constructing a system for monitoring potassium in patients with drug therapy that can adversely affect potassium. Antibiotics, dosing in the presence of renal failure, and digoxin dosing have been identified as additional topics of interest.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

This project represents a collaboration between faculty at Stanford University Medical Center and the University of Maryland School of Pharmacy in exploring computer-based monitoring of drug therapy. SUMEX, through its communications capabilities, facilitates this collaboration of geographically separated project participants by allowing development work on a central machine resource and file exchange between sites.

B. Sharing and Interactions with Other SUMEX-AIM Projects

Interactions with other SUMEX-AIM projects has been on an informal basis. Personal contacts have been made with individuals working on the ONCOCIN project concerning issues related to the formulation of the previously mentioned proposal. We expect interactions with other projects to increase significantly once the groundwork has been laid and issues directly related to AI are being addressed. Given the geographic separation of the investigators, the ability to exchange mail and programs via the SUMEX system as well as communicate with other SUMEX-AIM projects is vital to the success of the project.

C. Critique of Resource Management

To date, the resources of SUMEX have been fully adequate for the needs of this project. The staff have been most helpful with any problems we have had and we are quite satisfied with the current resource management. The only concerns we have relate to the state of the documentation on the system and the response time while using TYMNET from the Baltimore, Maryland area. While most aspects of the system are documented the path to a specific piece of information can be somewhat longer than one might expect. With respect to TYMNET, there are often up to 7 second pauses in the middle of transmissions. This can become quite annoying when trying to work with anything more than small bodies of text.

III. RESEARCH PLANS

A. Project Goals and Plans

The MENTOR project has the following goals:

1. Implement a prototype computer system to continuously monitor patient drug therapy in a hospital setting. This will be an expert system that will use a modular, frame-oriented form of medical knowledge, a separate inference engine for applying the knowledge to specific situations and automated collection of data from hospital information systems to produce therapeutic advisories.

2. Select a small number of important and frequently occurring medical settings (e.g., combination therapy with cardiac glycosides and diuretics) that can lead to therapeutic misadventures, construct a comprehensive medical knowledge base necessary to detect these situations using the information typically found in a computerized hospital information system and generate timely advisories intended to alter behavior and avoid preventable drug reactions.

3. Design and begin to implement an evaluation of the impact of the prototype MENTOR system on physicians' therapeutic decision-making as well as on outcome measures related to patient health and costs of care.

1985 will be spent on prototype development in four content areas, design and implementation of the basic knowledge representation and reasoning mechanisms and preliminary interfacing to existing patient information systems.

B. Justification and Requirements for Continued SUMEX Use

This project needs continued use of the SUMEX facilities for two reasons. First, it provides access to an environment specifically designed for the development of AI systems. The MENTOR project focuses on the development of such a system for drug monitoring that will explore some neglected aspects of AI in medicine. This environment is necessary for the timely development of a well-designed and efficient MENTOR system. Second, access to SUMEX is necessary to support the collaborative efforts of geographically separated development teams at Stanford and the University of Maryland.

The resources of SUMEX are central to the execution of the MENTOR project. A major component of the proposal was access to SUMEX resources and without it, the chances of funding would have been much less. Furthermore, the MENTOR project is predicated on the access to the SUMEX resource free of charge over the next two years. Given the current restrictions on funding, the scope of the project would have to be greatly reduced if there were charges for use of SUMEX.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

A major long-range goal of the MENTOR project is to implement this system on a independent hardware system of suitable architecture. It is recognized that the full monitoring system will require a large patient data base as well as a sizeable medical knowledge base and must operate on a close to real-time basis. Ultimately, the SUMEX facilities will not be suitable for these applications. Thus we intend to transport the prototype system to a dedicated hardware system that can fully support the the planned system and which can be integrated into the SUMC Hospital Information System. However, no firm decisions have been made about the requirements for this system since many specification and design decisions remain to be made.

D. Recommendations for Future Community and Resource Development

In the brief time we have been associated with SUMEX, we have been generally pleased with the facilities and services. However, it is clearly evident that the users almost insatiable demands for CPU cycles and disk space cannot be met by a single central machine. The best strategy would appear to be one of emphasizing powerful workstations or relatively small, multi-user machines linked together in a nation-wide network with SUMEX serving as the its central hub. This would give the individual users much more control over the resources available for their needs yet at the same time allow for the communications among users that have been one of SUMEX's strong points.

For such a network to be successful, further work needs to be done in improving the network capabilities of SUMEX to encourage users at sites other that Stanford. Specifically, the problem of slow throughput on TYMNET needs to be addressed for those users who do not have authorized access to ARPANET. Further work is also needed in the area of personal workstations to link them to such a network. Given the successful completion of this work, it would be reasonable to consider the gradual phase-out of the central SUMEX machine over two or three years to be replaced by an efficient, high-speed communications server.

IV.B.4. Rutgers Research Resource

Rutgers Research Resource--Artificial Intelligence in Medicine

Principal Investigators:
Casimir Kulikowski, Sholom Weiss
Rutgers University, New Brunswick, New Jersey

I. SUMMARY OF RESEARCH PROGRAM

A. Goals and Approach

The fundamental objective of the Rutgers Resource is to develop a computer based framework for advancing research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The central concept is to introduce advanced methods of computer science - particularly in artificial intelligence - into specific areas of biomedical inquiry. The computer is used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in problem solving and in processes of experimentation and theory formation.

An essential part of the resource is directed to methodological problems of knowledge representation and to the development of computer-based systems for acquiring, managing, and improving knowledge bases, and for constructing expert reasoning models in medicine. Equally fundamental are the problems of how to best use knowledge bases and models in processes of interpretation/diagnosis, planning, theory formation, simulation, and effective man-machine communication. These are problems we are studying in the Resource in the context of several system building efforts that address themselves to specific tasks of clinical decision-making and model development and testing.

Resource activities include research projects (collaborative research and core research) training/dissemination projects, and computing services in support of user projects.

B. Medical Relevance and Collaborations

In 1984-85 we continued the development of several versatile systems for building and testing consultation models in biomedicine. The EXPERT system has had many of its capabilities enhanced in the course of collaborative research in the areas of rheumatology, ophthalmology, and clinical pathology.

In *ophthalmology* we have developed a knowledge representation scheme for treatment planning which is both natural and efficient for encoding the strategies for choosing among competing and cooperating treatment plans. This involves a ranking of treatments according to their characteristics and desired effects as well as contraindications. A diagnosis and treatment planning program for ocular herpes was developed using this scheme. Our main collaboration continues to be with Dr. Chandler Dawson of the Proctor Foundation, UCSF.

In *rheumatology*, the model for rheumatological diseases now includes detailed diagnostic criteria for 26 major diseases. The management advice and treatment planning has been developed further. The Resource researchers have developed new representational elements for EXPERT in response to the needs of the rheumatology research Politakis originally developed a coordinated system called SEEK (System for

Empirical Experimentation with Expert Knowledge) which provides interactive assistance to the human expert in testing, refining and updating a knowledge base against a data base of trial cases. A generalized version of SEEK, SEEK2, has been developed during the past year. Dr. Lindberg of the National Library of Medicine, and Dr. Sharp, of the University of Missouri are the project leaders in developing the rheumatology knowledge base for this effort.

In *clinical pathology* our main collaboration has been with Dr. Robert Galen (Cleveland Clinic Foundation), with whom we have developed the serum protein electrophoresis model which is incorporated into an instrument a scanning densitometer. This instrument with interpretive reporting capabilities has now been on the market for over a year, is located at several hundred clinical sites. We are making good progress developing a knowledge based system for the interpretation of CPK/LDH isoenzymes.

In biomedical modeling applications we are experimenting with several prototype models for giving advice on the interpretation of experimental results in the field of enzyme kinetics, in conjunction with Dr. David Garfinkel. His PENNZYME program has been linked to a model in EXPERT, which allows the user to interpret the progress of the model analysis, and a framework for the design of experiments in this domain has been formulated.

C. Highlights of Research Progress

Research has continued on problems of representation, inference and control in expert systems. Emphasis has been placed this year on problems of knowledge base acquisition, empirical testing and refinement of reasoning (the SEEK2 system). From a technological point of view the market availability of the interpretive reporting version of a scanning densitometer, and the development of models for eye care consultation that run on microprocessor systems (Apple IIe, IBM-PC) represents an important achievement for AIM research in showing its practical impact in medical applications. This was recognized by the award of a scientific exhibit prize at the Academy of Ophthalmology Annual Meeting in November 1983.

- Knowledge Base Refinement: SEEK is a system which has been developed to give interactive advice about rule refinement during the design of an expert system. The advice takes the form of suggestions for possible experiments in generalizing and specializing rules in an expert model that has been specified based on reasoning rules cited by a human expert. Case experience, in the form of stored cases with known conclusions, is used to interactively guide the expert in refining the rules of a model. The design framework of SEEK consists of a tabular model for expressing expert-modeled rules and a general consultation system for applying a model to specific cases. This approach has proven particularly valuable in assisting the expert in domains where the logic for discriminating two diagnoses is difficult to specify; and we have benefited primarily from experience in building the consultation system in rheumatology. During the past year a newer SEEK2 system a been developed that has enhanced capabilities including a more generalized knowledge base and an automatic pilot capability to proceed with knowledge base refinements.

- Technology Transfer: Important technology transfer milestones have also been achieved this year: the instrument interpretation EXPERT program for serum protein has been widely disseminated as has the Ocular Herpes Treatment Program.

D. Up-to-Date List of Publications

The following is an update of publications in the Rutgers Resource for the period 1983 and 1984 (only publications not listed in previous SUMEX annual reports are presented here).

1. Apte, C. and Weiss, S.: *An Approach to Expert Control of Interactive Software Systems, IEEE Transactions on Pattern Analysis and Machine Intelligence* in press (1985).

2. Ginsberg, A., Weiss, S., and Politakis, P.: *SEEK2: A Generalized Approach to Automatic Knowledge Base Refinement* to appear in the Proceedings of the 1985 International Joint Conference on Artificial Intelligence.

3. Weiss, S.M. and Kulikowski, C.A.: *A Practical Guide to Designing Expert Systems*, Rowman and Allanheld, 1984.

4. Kastner, J., Weiss, S., Kulikowski, C., and Dawson, C.: *Therapy Selection in an Expert Medical Consultation System for Ocular Herpes Simplex* Computers in Biology and Medicine, Vol. 14, No. 3, pp. 285-301 (1984).

5. Dawson, C., Kastner, J., Weiss, S., Kulikowski, C.: *A Computer-based Method to Provide Subspecialist Expertise on the Management of Herpes Simplex Infections of the Eye*, Proceedings International Symposium On Herpetic Eye Diseases, Belgium (1984).

6. Galen, R. and Weiss, S.: *Predictive Value Calculator*, American Society of Clinical Pathologists, Clinical Chemistry # CC 84-4 (1984).

7. Kastner, J., Dawson, C., Weiss, S., Kern, K., Kulikowski, C.: *An Expert Consultation System for Frontline Health Workers in Primary Eye Care*, Journal of Medical Systems, Vol. 8, No. 5 (1984).

8. Kulikowski, C.A.: contributor to the Knowledge Acquisition chapter edited by B. Buchanan in the book *Building Expert Systems* (F. Hayes-Roth, et al., eds) Addison-Wesley, 1983.

9. Yao, Y. and Kulikowski, C.A.: *Multiple Strategies of Reasoning for Expert Systems*, Proc. Sixteenth Hawaii International Conference on Systems Sciences, pp. 510-514 , 1983.*

10. Kulikowski, C.A.: *Progress in Expert AI Medical Consultation Systems: 1980 - 1983* , Proc. MEDINFO '83 , pp. 499-502, Amsterdam, August 1983.*

11. Kastner, J.K., Weiss, S.M., and Kulikowski, C.A.: *An Efficient Scheme for Time-Dependent Consultation Systems*, Proc. MEDINFO '83, pp.619-622, 1983.*

12. Kulikowski, C.A.: *Expert Medical Consultation Systems*, Journal of Medical Systems, v.7, pp. 229-234, 1983.*

13. Weiss, S.M., Kulikowski, C.A., and Galen, R.S.: *Representing Expertise in a Computer Program: The Serum Protein Diagnostic Program*, Journal of Clinical Laboratory Automation, v.3, pp. 383-387, 1983.*

14. Kastner, J.K., Weiss, S.M., and Kulikowski, C.A.: *An Expert System for Front-line Health Workers in Primary Eye Care*, Proc. Seventeenth Hawaii International Conference on Systems Sciences, pp. 162-166, 1984.*

15. Kulikowski, C.A.: *Knowledge Acquisition and Learning in EXPERT*, Proc. 1983 Workshop on Machine Learning, Univ. of Illinois,Champaign-Urbana 1983.

Indicate by an asterisk (*) that the resource was given credit.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Dissemination

The SUMEX-AIM facility provides a backup node where some of our medical collaborators can access programs developed at Rutgers. The bulk of the medical collaborative work outlined in I.B. above is centered at the Rutgers facility (the Rutgers-AIM node).

Dissemination activities continue to be an important responsibility of the Rutgers Resource within the AIM community. The following activities took place in the last year:

1. Tenth AIM Workshop (1983):

   Organized by Dr. Chandrasekaran, it was held at Ohio State University. It consisted of a series of presentations on AIM research and related work by members of the AIM community.

2. 1984 Hawaii International Conference On Systems Sciences:

   Dr. Weiss presented a paper on the expert system for front-line health workers, and Dr. Kulikowski chaired a session on knowledge based medical systems.

B. National AIM Projects at Rutgers

The national AIM projects, approved by the AIM Executive Committee, that are associated with the Rutgers-AIM node are the following:

1. INTERNIST/CADUCEUS project, headed by Dr. Myers and Dr. Pople from the University of Pittsburgh, has been using the Rutgers Resource as a backup system for development and experimentation.

2. Medical Knowledge Representation project, headed by Dr. Chandrasekaran from Ohio State University, is doing most of its research on the Rutgers system.

3. PURSUIT project, directed by Dr. Greenes from Harvard University, is doing most of its research on a Goal-Directed Model of Clinical Decision-Making at Rutgers.

4. Biomedical Modeling, by Dr. Garfinkel from the University of Pennsylvania.

5. Attending Project, directed by Dr. Perry Miller of the Yale Medical Center, is doing much of the research on critiquing a physician's plan of management at Rutgers.

6. MEDSIM project: This is a pilot project designed to provide resource-sharing and community building facilities for about 25 researchers in bio-mathematical modeling and simulation.

C. Critique of SUMEX-AIM Resource Management

Rutgers is currently using the SUMEX DEC-20 system primarily for communication with other researchers in the AIM community and with SUMEX staff, and also for backup computing in demonstrations, conferences and site visits. Our usage is currently running at less than 50 connect hours per year at SUMEX, with an overall connect/CPU ratio of about 30.

Rutgers is beginning to place more emphasis on the use of personal computers, and on network support needed to make these effective. SUMEX has been of significant help in their developmental efforts in networking workstation software.

III. RESEARCH PLANS

A. Project Goals and Plans

We are planning to continue along the main lines of research that we have established in the Resource to date. Our medical collaborations will continue with emphasis on development of expert consultation systems in rheumatology, ophthalmology and clinical pathology. The basic AI issues of representation, inference and planning will continue to receive attention. Our core work will continue with emphasis on further development of the EXPERT framework and also on AI studies in representations and problems of knowledge and expertise acquisition. We propose to work on a number of technology transfer experiments to micro processing that will be affordable by our biomedical research and clinical collaborators. We also plan to continue our participation in AIM dissemination and training activities as well as our contribution -- via the Rutgers computers -- to the shared computing facilities of the national AIM network.

B. Justification and Requirements for Continued SUMEX Use

Continued access to SUMEX is needed for:

1. Backup for demos, etc.

2. Programs developed to serve the National AIM Community should be runnable on both facilities.

3. There should be joint development activities between the staffs at Rutgers and SUMEX in order to ensure portability, share the load, and provide a wider variety of inputs for developments.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

Our computing needs are based on a centralized computing resource accessible to distant users, and local workstations. We will continue to use Sumex for backup purposes.

D. Recommendations for Future Community and Resource Development

Use of personal computers and workstations is continuing to grow in the AIM community. We find that the biggest challenge is supporting these systems. Although some central computing will continue to be needed for communication and coordination, we believe that over the next few years all AIM research projects and even individual collaborators will come to have their own hardware. However many of these community members (particularly the collaborators) will not be in a position to support hardware or software on their own. We would certainly expect SUMEX to continue to provide expert advice in this area. However we believe it would be helpful for SUMEX to have a formal program to support smaller computers in the field. We envision this as including at least the following items:

- A central source of information on hardware and software that is likely to be of interest to the AIM community. SUMEX might want to become a distribution point for certain of this software, and even help coordinate quantity purchase of hardware if this proves useful.

- Assistance in support of hardware and software in the field. Depending upon the hardware involved, this might involve advice over the telephone or actual board-swapping by mail.

IV.B.5. SECS: Simulation & Evaluation of Chemical Synthesis

SECS - Simulation and Evaluation of Chemical Synthesis Project

Principal Investigator: W. Todd Wipke
Board of Studies in Chemistry
University of California
Santa Cruz, CA. 95064

Coworkers:

I. Kim	(Grad student)
M. Hahn	(Grad Student)
M. Yanaka	(Postdoctoral)
I. Iwataki	(Postdoctoral)
T. Okada	(Postdoctoral)

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

With the SECS project our long range goal is to develop the logical principles of molecular construction and to use these in developing practical computer programs to assist investigators in designing stereospecific syntheses of complex bio-organic molecules. Our second area of research, the XENO project, is aimed at improving methods for predicting potential biological activity of metabolites and plausibility of incorporation and excretion of metabolites.

B. Medical Relevance and Collaboration

The development of new drugs and the study of drug structure biological activity relationships depends upon the chemist's ability to synthesize new molecules as well as his ability to modify existing structures, e.g., incorporating isotopic labels or other substituents into bio-molecular substrates. The Simulation and Evaluation of Chemical Synthesis (SECS) project aims at assisting the synthetic chemist in designing stereospecific syntheses of biologically important molecules. The advantages of this computer approach over normal manual approaches are many: 1) greater speed in designing a synthesis; 2) freedom from bias of past experience and past solutions; 3) thorough consideration of all possible syntheses using a more extensive library of chemical reactions than any individual person can remember; 4) greater capability of the computer to deal with the many structures which result; and 5) capability of computer to see molecules in a graph theoretical sense, free from the bias of 2-D projection.

The objective of using XENO in metabolism studies is to predict the plausible metabolites of a given xenobiotic in order that they may be analyzed for possible carcinogenicity. Metabolism research may also find this useful in the identification of metabolites in that it suggests what to look for. Finally, one may envision applications of this technology in problem domains where one wishes to alter molecules in order to inhibit certain types of metabolism.